

ABELIAN GIRTH AND GIRTH

JOEL FRIEDMAN, ALICE IZSAK, AND LIOR SILBERMAN

ABSTRACT. We show that the abelian girth of a graph is at least three times its girth. We prove an analogue of the Moore bound for the abelian girth of regular graphs, where the degree of the graph is fixed and the number of vertices is large. We conclude that one could try to improve the Moore bound for graphs of fixed degree and many vertices by trying to improve its analogue concerning the abelian girth.

CONTENTS

1. Introduction	1
2. Main Results	4
2.1. Graph Terminology	4
2.2. Our Fundamental Lemma	5
2.3. Main Results in this Paper	7
3. Proof of Theorem 2.8	7
4. Proofs of the Fundamental Lemma and Theorem 2.7	10
5. Abelian Girth of the LPS Expanders	12
References	14

1. INTRODUCTION

In the work [Fri15] the first named author introduced the *abelian girth* of a graph, a graph invariant, and used it in a proof of the Hanna Neumann Conjecture. In this paper we study the abelian girth and its relationship to the ordinary girth and to the volume bound on girth known as the Moore bound. Specifically, we give some evidence to show that bounding the abelian girth from above is a possible approach to improving the Moore bound for regular graphs of fixed degree with a large number of vertices. Let us make this more precise.

2010 *Mathematics Subject Classification.* 05C38.

Key words and phrases. graph theory, girth.

Research of all authors partly supported by an NSERC Discovery Grants.

The *girth* of a graph is the length of the shortest non-trivial cycle in the graph. If $G_{n,d}$ is any d -regular graph on n vertices, it is known that for fixed d and large n we have

$$(1) \quad \text{Girth}(G_{n,d}) \leq 2 \log_{d-1} n + o_n(1),$$

and we are interested to know if the factor of 2 can be improved upon. This bound follows from the *Moore bound* (see [ERS66].) Although the Moore bound for regular graphs is very easy to prove, there has been only slight improvements to it in the last 50 years—only an additive constant of one or two. In particular, the multiplicative factor of 2 in (1) is the best factor known to date. Graphs of large girth have numerous applications ([ABLS07, TZ05, MN96]) and much has been written on the girth and the Moore bound (see [AHL02]).

The *abelian girth* of a graph, G , denoted $\text{Abl}(G)$, is the girth of its *universal abelian cover*, or equivalently the shortest length of a non-trivial, closed, non-backtracking walk that traverses each edge the same number of times in each direction. The abelian girth is important in sheaf theory on graphs and the first proof of the Hanna Neumann Conjecture [Fri15]; furthermore, in a forthcoming paper we will show the abelian girth is strongly related to what we call “gapped” sheaves on a graph [FIS]. In this paper we show that there is an analogue of the Moore bound for the abelian girth, and improving this analogue would improve the Moore bound. Furthermore, we provide evidence that such an improvement may be possible.

Specifically, we show that

$$(2) \quad \text{girth}(G) \leq 3 \text{Abl}(G),$$

for any graph, G . Second, we show that there is an argument analogous to the Moore bound that shows that for fixed d we have

$$(3) \quad \text{Abl}(G_{n,d}) \leq 6 \log_{d-1} n + o_n(1)$$

for any d -regular graph on n vertices, $G_{n,d}$; we remark that the proof of this theorem is not as immediate as that of the Moore bound. It follows that any improvement to the factor of 6 in (3) would give an improvement to the factor of 2 in (1).

Ideally we would show that all known explicit constructions of families of d -regular graphs have abelian girth at most $c \log_{d-1} n$ for some $c < 6$. In this paper, we will focus on the only family of d -regular graphs with d fixed which has girth greater than $\log_{d-1} n$. It is known that the factor of 2 in (1) cannot be less than $4/3$, at least for certain d : indeed, [LPS88] constructs graphs, $X^{p,q}$, for primes $p, q \equiv 1 \pmod{4}$,

that are $d = p + 1$ regular on $n = q(q^2 + 1)$ vertices for which

$$\text{girth}(X^{p,q}) = (4/3) \log_{d-1} n + o_n(1)$$

for fixed $d = p + 1$ and large $n = q(q^2 + 1)$. Margulis [Mar88] independently constructed similar graphs with girth asymptotically as large as that of the graphs from [LPS88]. Furthermore there are no known families of graphs which improve on the above $4/3$; in fact, for general d , the best girth lower bound is $4/3$ replaced with 1, by choosing a random graph and slightly modifying it [EL14], [ES63]. To show that it is plausible to improve on the factor of 6 in (3), we will show that

$$(4) \quad \text{Abl}(X^{p,q}) \leq (16/3) \log_{d-1} n + o_n(1)$$

which suggests that there is room for the factor of 6 to decrease. We conjecture that the $16/3$ can be replaced with 4, for reasons we shall explain later. We are unaware of any other graph constructions in the literature which improves upon the $16/3$ above.

To prove (2), we prove a fundamental lemma that suggests many possible generalizations of the above discussion of abelian girth and the Moore bound. First we remark that the girth of a graph is smallest positive length of a cycle that embeds in the graph. Our fundamental lemma states that the abelian girth is the smallest number of edges in a graph of Euler characteristic -1 that embeds in the graph, provided that we weight the edges appropriately: namely, we count each edge twice, except that in “barbell graphs” we count each edge in the bar four times.

We remark that both girth and abelian girth can be viewed as linear algebraic invariants of graphs; indeed, girth is often studied as a property of the adjacency matrix of the graph (see, for example, [AHL02]), and the abelian girth relates to sheaves of vector spaces over the graph. So our lemma fundamental to proving (2) also suggests that there could be other such girth-type invariants, both (1) arising as linear algebraically from the graphs, and (2) satisfying inequalities analogous to (2).

The rest of this paper is organized as follows. In Section 2 we give some precise definitions and state our main theorems. In Section 3 we prove (3). In Section 4 we prove the fundamental lemma as well as (2) which follows quickly from the fundamental lemma. In Section 5 we describe the graphs $X^{p,q}$ from [LPS88] and show that they obey Equation (4).

2. MAIN RESULTS

In this section we fix some terminology regarding graphs and formally state our main results.

2.1. Graph Terminology. This entire subsection consists of definitions used throughout this paper; they are more or less standard.

By a *graph* we shall mean a quadruple $G = (V_G, E_G, t_G, h_G)$ where V_G and E_G are sets (the “*vertices*” and “*edges*”, respectively) and $t_G, h_G: E_G \rightarrow V_G$ are maps (the “tail” and “head” of each edge, respectively). The *Euler characteristic* of G is

$$\chi(G) = |V_G| - |E_G|,$$

where $|\cdot|$ denotes the cardinality, provided that G is finite, i.e., that $|V_G|$ and $|E_G|$ are finite. We define the *directed edge set* of G to be

$$D_G = E_G \times \{+, -\},$$

and we extend h_G and t_G to be functions on D_G such that for $e \in E_G$

$$h_G(e, +) = t_G(e, -) = h_G(e), \quad t_G(e, +) = h_G(e, -) = t_G(e);$$

for $e \in E_G$ we say that $(e, +)$ and $(e, -)$ are *inverses* of each other. We denote the inverse of $u \in D_G$ by u^{-1} . A *walk* of *length* m in G is a sequence of directed edges

$$w = (u_1, \dots, u_m)$$

such that the head of u_i is the tail of u_{i+1} for $i = 1, \dots, m-1$; we define the vertices $t_G(u_1)$ and $h_G(u_m)$ to be the *starting* and *terminating* vertices of w , jointly the *endpoints* of w , and the *length* of w to be m , denoted $l(w)$; we say that w is *closed* if its two endpoints are the same vertex; we refer to the *edges* of w as the $e \in E_G$ such that at least one of $(e, +), (e, -)$ is an directed edge of w ; we say that w is *non-backtracking* if there is no $i = 1, \dots, m-1$ for which $u_i^{-1} = u_{i+1}$; we say that w is *strongly closed, non-backtracking* if w is closed and non-backtracking, and u_1, u_m are not inverses of each other; we say that w is a *path* (respectively, *cycle*) if w is non-backtracking with distinct endpoints (respectively strongly closed, non-backtracking), and each edge $e \in E_G$ appears at most once in w (i.e. $(e, +)$ and $(e, -)$ appears at most once in u_1, \dots, u_m). The *inverse* of w is defined as the walk

$$w^{-1} = (u_m^{-1}, \dots, u_1^{-1}).$$

Given a walk $w = (u_1, \dots, u_m)$ and a walk $k = (t_1, \dots, t_n)$ such that the terminating vertex of w is the starting vertex of k , we define the *product* wk to be the walk $(u_1, \dots, u_m, t_1, \dots, t_n)$. We say that a walk

w as above *joins* its two endpoints. We say that G is *connected* if any two of its vertices are joined by some walk in G .

The contiguous appearance of an edge and its inverse in a walk is called a *reversal*. Each walk, w , can be *reduced* by successively discarding its reversals; the walk obtained is the *reduction* of w and is known to be independent of the choice of pairs to reduce; the reduction therefore contains no reversals and is non-backtracking [DD89]. We use the notation $\text{red}(w)$ for the reduction of w . Note that $\text{red}(w^{-1}) = \text{red}(w)^{-1}$ for any walk w since the portion of a walk that are removed by reduction remain the same when we take the inverse of that walk.

Definition 2.1. In each element ω of $\pi_1(G, u)$, there exists one unique non-backtracking walk in G ; any walk in ω reduces to that unique non-backtracking walk [DD89]. By the *length* of ω , denoted $|\omega|$, we mean the length of the unique non-backtracking walk in ω . Note this is not word length with respect to a set of generators of $\pi_1(G, u)$. Similarly, even though closed walks are members of homotopy classes in the fundamental group, when we refer to the length of a walk or the reduction of a walk w we mean length and reduction as defined earlier and not in with respect to a set of generators for a free group.

As usual, a *morphism* of graphs, $f: G \rightarrow H$, is a pair $f = (f_V, f_E)$ of maps $f_V: V_G \rightarrow V_H$ and $f_E: E_G \rightarrow E_H$ such that $t_H \circ f_E = f_V \circ t_G$ and $h_H \circ f_E = f_V \circ h_G$. We often drop the subscripts from f_V and f_E . The set of morphisms will be denoted $\text{Mor}(G, H)$. Thus morphisms are homomorphisms of the underlying undirected graphs which preserve the orientation.

As usual, the *girth* of a graph is the length of its shortest closed, non-backtracking walk; this length is necessarily positive by our conventions above. (In the literature one often allows for walks of length zero, which we do not consider here.)

2.2. Our Fundamental Lemma. In this subsection we discuss our main results.

Definition 2.2. The *abelian girth* of a graph, G , denoted $\text{Abl}(G)$, is the minimum $m \geq 1$ such that there is a closed, non-backtracking walk

$$w = (u_1, \dots, u_m)$$

such that each edge is traversed the same number of times in w in both directions, i.e., for each $e \in E_G$ the edge $(e, +)$ appears the same number of times among u_1, \dots, u_m as does $(e, -)$.

The abelian girth is the same as the girth of the *universal abelian cover* of G ; see [Fri15]. Given a walk w and an $e \in E_G$, if $(e, +)$ appears

i_+ times and $(e, -)$ appears i_- times we say e appears a net $i_+ - i_-$ times in w . We refer to a walk as *edge neutral* if every edge appears a net 0 times in it.

We begin with a fundamental lemma. We remind the reader that our conventions insists that cycles and paths are of positive length.

Definition 2.3. Let G be a connected graph of Euler characteristic -1 without leaves, i.e., without vertices of degree one. We say that G is

- (1) a *figure-eight* graph if G consists of two cycles, mutually edge disjoint, sharing the same endpoint;
- (2) a *barbell* graph if G consists of two cycles, w_1, w_2 , and one path b , all mutually edge disjoint, such that b joins the endpoints of w_1 and w_2 ; we refer to b as the *bar* of G ;
- (3) a *theta* graph if G consists of two vertices joined by three paths mutually edge disjoint.

It is well known that the above three cases classifies all connected graphs of Euler characteristic -1 without leaves (see, for example, [LP10].)

Definition 2.4. Let G be a connected graph of Euler characteristic -1 without leaves. We define the *abelian length* of G , denoted $l_{\text{Abl}}(G)$, to be twice its number of edges except that each edge of its bar (if G is a figure-eight graph) is counted four times.

Remark 2.5. Observe that in each case, the abelian length of such G is its abelian girth. This is clear for the figure-eight and the theta (traverse each cycle twice) and not hard for the barbell (must traverse each cycle twice, in reverse orientations, so must traverse the bar between each of the four cycle traversals).

Lemma 2.6. (The Fundamental Lemma) *For any graph, G , we have*

$$\text{Abl}(G) = \min_{G' \subset G} l_{\text{Abl}}(G')$$

taken over all subgraphs, G' , that are connected, of Euler characteristic -1 , and without leaves. If no such G' exist, then the above minimum is taken to be infinity, as is the abelian girth of G , and there are no closed non-backtracking walks that traverse each edge the same number of times in both directions.

By the observation above, the upper bound is clear, and the content is in the lower bound.

2.3. Main Results in this Paper. Now we can easily state the other main results in this paper.

Theorem 2.7. *For any graph we have*

$$\text{Girth}(G) \leq \text{Abl}(G)/3.$$

This theorem is an easy corollary of Lemma 2.6.

Theorem 2.8. *For any fixed d , let $G_{n,d}$ be any d -regular graph on n vertices. Then for large n we have*

$$\text{Abl}(G) \leq 6 \log_{d-1} n + o_n(1).$$

Our last theorem uses the [LPS88] graphs $X^{p,q}$ whose definition we save for Section 5.

Theorem 2.9. *Let $p, q \equiv 1 \pmod{4}$ be prime with $(q/p) = 1$. The graph $X^{p,q}$ of [LPS88] of our Definition 5.1 has degree $d = p + 1$ and $n = q(q^2 + 1)$, and for fixed p and large q (i.e., fixed d and large n) we have*

$$\text{Abl}(X^{p,q}) \leq (16/3) \log_{d-1} n + o_n(1).$$

3. PROOF OF THEOREM 2.8

In this section we prove Theorem 2.8. First we state a few well-known facts regarding free groups and graphs.

Lemma 3.1. *Let G be a finite graph. Then the fundamental group, $\pi_1(G, u)$, of homotopy classes of closed walks in G about u for any $u \in V_G$ is isomorphic a free group on a finite set of generators.*

The following fact is well known. For ease of reading we provide a proof which is also found in texts on group theory such as [LS15].

Lemma 3.2. *Let F be a free group and let α and β be two elements of F that commute with each other. Then there exists an $\omega \in F$ such that $\alpha = \omega^m$ and $\beta = \omega^{m'}$ for $m, m' \in \mathbb{Z}$.*

Proof. By the Nielsen-Schreier theorem, the subgroup that α, β generate is a free group. And yet, any two elements of this subgroup commute, and hence this subgroup cannot be free and rank more than one. Hence this subgroup is generated by some $\omega \in F$, and the conclusion follows. \square

For any closed walk w , we use the notation $[w]$ for the homotopy class in $\pi_1(G, u)$ that has w as a representative. Suppose a and b are closed walks and $[a]$ and $[b]$ commute in $\pi_1(G, u)$. Note that equality in the

previous lemma is with regards to $\pi_1(G, u)$, meaning an equality of homotopy classes but not necessarily walks. Reducing all the elements of a homotopy class gives one unique walk though. So since $[a]$ and $[b]$ commute, we have $\text{red}(a) = \text{red}(w^m)$ and $\text{red}(b) = \text{red}(w^{m'})$ for $m, m' \in \mathbb{Z}$ and w a closed walk. These equalities may not exist without the reductions. We may assume w is reduced, since $\text{red}(w^m) = \text{red}(\text{red}(w)^m)$.

The following lemma is also well known and easy.

Lemma 3.3. *Let w be a closed nontrivial non-backtracking walk in a graph. Then if for integers m, m' we have that the length of $\text{red}(w^m)$ equals that of $\text{red}(w^{m'})$, then $m = \pm m'$.*

Proof. Let y be a maximal length walk such that $w = yxy^{-1}$ for some walk x . Then it is easy to check that for $m \neq 0$ the length of $\text{red}(w^m)$ with respect to S is precisely

$$2|y| + |m||x|.$$

□

Proof of Theorem 2.8. Fix any vertex, $v \in V_G$. Then for any integer $h \geq 1$, there are $d(d-1)^{h-1}$ non-backtracking walks of length h from v . So let h be the smallest integer for which

$$d(d-1)^{h-1} \geq 2n+1;$$

then, by the pigeon hole principle, there are three distinct non-backtracking walks, a, b, c , in G beginning in v and terminating in the same vertex u (which may or may not equal v). We remark that

$$h = \log_{d-1} n + o_n(1).$$

Hence, to prove the theorem it suffices to show that we have

$$(5) \quad \text{Abl}(G) \leq 6h.$$

The rough idea is simple. We break the analysis into two cases: $u = v$ and $u \neq v$. In either case we define new walk, ℓ , based on a, b, c such that (1) the length of ℓ is $4h$ or $6h$ (in the respective two cases), (2) each edge of E_G appears a net 0 times in ℓ , but (3) ℓ is not necessarily non-backtracking. The main work is to show that ℓ does not reduce to the empty word. Discarding a consecutive pair of an (unoriented) edge of ℓ and its opposite retains the property that each edge is traversed the same number of times in each direction. Hence if ℓ is reduced to a non-empty, non-backtracking walk, ℓ' , then ℓ' is edge neutral, and so

$$\text{Abl}(G) \leq 6h.$$

Let us describe ℓ as above. If $u = v$, then at least two of a, b, c are not inverses of each other; if a, b are such walks, then we set $\ell = aba^{-1}b^{-1}$.

This walk is of length $4h$. If $u \neq v$, then we set ℓ to $\ell = ab^{-1}ca^{-1}bc^{-1}$, which is a walk of length $6h$. For any $e \in E_G$, if e appears a net k times in a walk w then e appears a net $-k$ times in the walk w^{-1} . Thus ℓ in either case is edge neutral.

It remains to show that ℓ reduces to a non-empty non-backtracking walk.

The case $u = v$ and $\ell = aba^{-1}b^{-1}$ where a and b (are distinct and) are not inverses of each other is relatively easy. If ℓ reduces to the empty word e and if $\alpha = [a]$, $\beta = [b]$ and $\delta = [\ell]$ then $\delta = [e] = [\alpha, \beta]_{\text{comm}}$ where $[\ , \]_{\text{comm}}$ denotes the commutator. So α and β commute as elements of $\pi_1(G, u)$. Hence, by Lemma 3.2, and since $\pi_1(G, u)$ is a free group, it follows that $a = \text{red}(w^m)$ and $b = \text{red}(w^{m'})$ for some closed walk w with $w \neq 1$; but then by Lemma 3.3, since a and b have the same length, we have $m = \pm m'$, which contradicts the fact that a and b are distinct and not inverses of each other.

Now suppose $u \neq v$. We remark that as elements of $\pi_1(G, u)$ we have

$$\ell = ab^{-1}cb^{-1}ba^{-1}bc^{-1} = [ab^{-1}, cb^{-1}],$$

where we use $[f_1, f_2]$ to refer to the walk $f_1 f_2 f_1^{-1} f_2^{-1}$ for any f_1 and f_2 that are closed walks about the same vertex. Let $\delta = [\ell]$, $\alpha = [ab^{-1}]$ and $\gamma = [cb^{-1}]$ and let e be the trivial walk. So if $\ell \in [e]$, we have that $1 = \delta = [\alpha, \gamma]_{\text{comm}}$ and so $[ab^{-1}]$ and $[cb^{-1}]$ commute in the fundamental group. Hence $\text{red}(ab^{-1}) = \text{red}(w^m)$ and $\text{red}(cb^{-1}) = \text{red}(w^{m'})$ for some closed non-backtracking walk w . Let $w = yxy^{-1}$ with x a closed non-backtracking walk and y a walk of maximal length; then (1) $|x| > 0$ as $|x| = 0$ would imply $a = b$, (2) the first and last edges of x are not inverses of each other, and (3) $\text{red}(w^m) = yx^m y^{-1}$.

Next we want to make some remarks on ab^{-1} and w^m based on the fact that $\text{red}(ab^{-1}) = \text{red}(w^m)$; we will later repeat similar remarks for cb^{-1} and $w^{m'}$ based on the fact that $\text{red}(cb^{-1}) = \text{red}(w^{m'})$. Note

$$a = \text{red}(ab^{-1}b) = \text{red}(\text{red}(ab^{-1})b) = \text{red}(\text{red}(w^m)b) = \text{red}(w^m b)$$

as reductions are independent of the order in which we reduce portions of the word.

Let p be the maximal prefix of b whose inverse is a suffix of $\text{red}(w^m) = yx^m y^{-1}$ (a priori p could be as long as the shorter of $yx y^{-1}$ and b); since

$$a = \text{red}(w^m b),$$

we have

$$|a| = 2|y| + |m||x| + |b| - 2|p|.$$

But by assumption $|a| = |b|$, and hence

$$2|y| + |m||x| = 2|p|.$$

It follows that

$$(6) \quad |p| = |y| + |m| |x|/2 > |y|,$$

since $|x| > 0$ and $m \neq 0$. Since $|p| > |y|$ and p is a suffix of $w^m = yx^m y^{-1}$, it follows that p begins with y and contains at least one more edge; hence $p = yx_1$ where $|x_1| > 0$ and the first edge of x_1 is the inverse of the last edge of x^m .

However, since $\text{red}(cb^{-1}) = \text{red}(w^{m'})$, the very same arguments show that if p' is defined analogously, i.e., as the maximal prefix of b that is a suffix of $\text{red}(w^{m'})$, then $p' = yx'_1$ where $|x'_1| > 0$ and the first letter of x'_1 is the inverse of the last letter of $x^{m'}$. But since both p and p' are both prefixes of b , we must have that m and m' have the same sign: for if, say, $m > 0$ and $m' < 0$, then the last edge of x^m or its orientation is different than the last edge of $x^{m'}$ by the maximality of y in the equation $w = yxy^{-1}$.

Hence

$$a = \text{red}(w^m b), \quad c = \text{red}(w^{m'} b),$$

a, b, c have the same length, and without loss of generality we may assume $m > m' > 0$. But then (6) and its analogue for $cb^{-1} = w^{m'}$ show that

$$m = (|p| - |y|)/|x|, \quad \text{and} \quad m' = (|p'| - |y|)/|x|.$$

It follows that $|p| > |p'|$. Since both p and p' are prefixes of b , we have that p' is a prefix of p . But p' is the maximal suffix of $w^{m'}$ that is also a prefix of b ; since $w^{m'}$ is also a suffix of w^m and w^m is a prefix of b it follows that $p' = w^{m'}$ (otherwise p could not be larger than p'). But in this case

$$|c| = |b| - |p'| < |b|,$$

which contradicts that fact that $|c| = |b|$ from our choice of a, b and c . \square

Lemma 3.4. *If a graph G has three distinct closed, non-backtracking walks from a vertex to itself and each walk has length l then $\text{Abl}(G) \leq 4l$.*

Proof. This was shown to be true in our proof of Theorem 2.8. See the discussion for the case $u = v$. \square

4. PROOFS OF THE FUNDAMENTAL LEMMA AND THEOREM 2.7

Proof of Lemma 2.6. Let w be an edge neutral, closed, non-backtracking walk of minimal (positive) length, and let B be the subgraph of G of vertices and edges that occur in w . Then B is a connected subgraph of G such that each vertex has degree at least two. If every

vertex in B has degree exactly two, then B would be a cycle, which is impossible any non-backtracking walk in a cycle traverses edges in at most one direction. Hence at least some vertex of B is strictly greater than two; hence the formula

$$\chi(B) = \sum_{v \in V_B} (2 - \deg(v))$$

shows that $\chi(B) \geq -1$.

Now w traverses each edge of B at least once in each direction, and hence

$$l(w) \geq 2|E_B|.$$

Recall that the abelian length of a figure-eight, theta or barbell graph was defined as its abelian girth. In particular, if B is such a graph we are done. Moreover, in general it suffices to show that there is another walk w' which is edge neutral and closed non-backtracking for which $l(w) \geq l(w')$ and which is supported on a subgraph B' of B of Euler characteristic -1 . We therefore assume by contradiction that no such w' exists.

It follows immediately that B cannot contain a theta or figure-eight graph B' ; indeed

$$l(w) \geq 2|E_B|;$$

while for such B' , $Abl(B') = 2|E_{B'}|$ and B' would be a proper subgraph of B , hence have fewer edges. This would contradict the minimality of B .

Second, we claim that B cannot be 1-connected (have an edge e whose removal would disconnected B). Indeed, if there is such an edge, we can extend in both directions to a path until we encounter vertices v_1, v_2 of degree greater than 2. Call this path p .

Let B_1, B_2 be the components we get by removing p from B (but keeping $v_1 \in B_1$ and $v_2 \in B_2$). Then any non-backtracking closed walk in B must follow a traversal of p from B_1 to B_2 by a non-trivial closed walk in B_2 , a traversal of p , and then a non-trivial walk in B_1 . In particular, each of B_1, B_2 must contain a cycle and w must traverse p at least 4 times, between successive visit to B_1, B_2, B_1, B_2 .

Each visit to B_i is a closed non-backtracking walk in B_i beginning in v_i . Accordingly let c_i be the shortest such walk. Then we may assume w is the sequence $c_1, p, c_2, p^{-1}, c_1^{-1}, p, c_2^{-1}, p$ (this is no longer than w), and in particular that B_i is the support of a shortest closed cycle. But then B_i is unicyclic, and since it has no leaves it must be a "lollipop": a simple cycle connected to v_i by a "stick" (simple path). This which

makes B into a barbell graph and w into its shortest balanced closed non-backtracking walk.

Third, we claim that either the first or second claim must be contradicted; this will complete the proof of the lemma. Indeed, B is not a tree, and therefore must contain a cycle, C . The cycle must contain a vertex, v , of degree at least three, or else B would consist entirely of C , which would then have Euler characteristic zero. Let e be any edge not in C that is incident upon v , and let u be the other endpoint of e . Since removing e does not disconnect B , there must be a walk from u to a vertex of C that does not contain e ; let p be such a walk of minimal length. Then p begins in u and is a beaded path to a vertex of c . But then the path e followed by p is disjoint from C and joins v to another vertex of C , which yields either a figure-eight graph (if p terminates in v) or a theta graph (if p terminates in a vertex of C that is not v).

Hence our first claim, that B does not contain a theta or figure-eight subgraph, is violated. \square

We now prove Theorem 2.7 as a consequence of the Fundamental Lemma.

Proof of Theorem 2.7. Let $G' \in \mathcal{C}$ be the subgraph of G of minimal abelian length. If G' is homeomorphic to the barbell graph, then $\text{Abl}(G) \geq 4 + 4 \text{Girth}(G)$ since each of the two simple cycles in G' is at least as long as the girth. Similarly, $\text{Abl}(G) \geq 4 \text{Girth}(G)$ if G' is homeomorphic to a figure-eight graph. In the case of a theta graph, we have two vertices u and v in G' of degree 3 and three simple paths from u to v . Let the lengths of the simple path be l , m and n . Then $l + m, m + n, n + l \geq \text{Girth}(G)$. Adding these inequalities gives

$$3 \text{Girth}(G) \leq 2(l + m + n) = 2|E_{G'}| = \text{Abl}(G)$$

. \square

5. ABELIAN GIRTH OF THE LPS EXPANDERS

The Ramanujan graphs of Lubotzky, Phillips and Sarnak [LPS88] are an infinite family of graphs with the largest known asymptotic girth. As mentioned earlier, these are graphs $X^{p,q}$, for primes $p, q \equiv 1 \pmod{4}$, that are $d = p + 1$ regular on $n = q(q^2 + 1)$ vertices for which

$$\text{girth}(X^{p,q}) = (4/3) \log_{d-1} n + o_n(1).$$

This immediately leads to the bound on abelian girth

$$\text{Abl}(X^{p,q}) \geq 4 \log_{d-1} n + o_n(1).$$

We will describe these Ramanujan graphs and obtain an upper bound on their abelian girth in order to show that their abelian girth isn't so large that it would be impossible to improve Theorem 2.8.

Definition 5.1. Let p and q be unequal primes congruent to 1 mod 4 with $(p/q) = -1$ and $q > \sqrt{p}$. The integral quaternions, denoted $H(\mathbb{Z})$, are given by

$$H(\mathbb{Z}) = \{\alpha = \alpha_0 + \alpha_1\mathbf{i} + \alpha_2\mathbf{j} + \alpha_3\mathbf{k} \mid \alpha_j \in \mathbb{Z}\}.$$

We denote the conjugate of α by $\bar{\alpha}$ and define $N(\alpha) = \alpha\bar{\alpha}$. Let S be the set of all α in $H(\mathbb{Z})$ satisfying $N(\alpha) = p, \alpha \equiv 1(\text{mod } 2)$ and $\alpha_0 \geq 0$. It can be shown that $|S| = p + 1$. Define $\Lambda'(2)$ as the set of $\alpha \in H(\mathbb{Z})$ such that $N(\alpha) = p^v$ for some non-negative integer v and $\alpha \equiv 1(\text{mod } 2)$. Define $\Lambda(2)$ as equivalence classes that identify α and $\beta \in H(\mathbb{Z})$ if $\pm p_1^v \alpha = p_2^v \beta$ for some $v_1, v_2 \in \mathbb{Z}$. It is known that the Cayley graph of $\Lambda(2)$ is the $(p + 1)$ -regular tree. Define $\Lambda(2q)$ by

$$\Lambda(2q) = \{[\alpha] \in \Lambda(2) \mid 2q \text{ divides } \alpha_j, j = 1, 2, 3\}.$$

This is a normal subgroup of $\Lambda(2)$ and $X^{p,q}$, the *LPS graph*, is defined as the Cayley graph of $\Lambda(2)/\Lambda(2q)$ with generators $S/\Lambda(2q)$. This graph is known to be a $(p + 1)$ -regular bipartite graph on $q(q^2 + 1)$ vertices.

Since $X^{p,q}$ is a Cayley graph it is vertex transitive which allows us to assume its smallest cycle goes from the identity element of $\Lambda(2)$ to some nontrivial element of $\Lambda(2q)$ along the infinite Cayley graph of $\Lambda(2)$. We now define a vertex' depth as its distance from the identity in the Cayley graph of $\Lambda(2)$. Biggs and Boshier [BB90] show that if $[b] \in \Lambda(2)$ is at depth $2r$ and $r > 0$ then there is some b in the equivalence class $[b]$ such that

$$b_0 = \pm(p^r - mq^2)$$

with $m > 0$ and even.

Definition 5.2. Positive integers are called *good* if they are not of the form $4^\alpha(8\beta + 7)$ for integers $\alpha, \beta \geq 0$.

The following is Lemma 2 of [BB90].

Lemma 5.3. *There exists a $[b] \in \Lambda(2q)$ at level $2r$ with $b_0 = p^r - mq^2$ with $m > 0$ and b_0 positive if and only if $2mp^r - m^2q^2$ is good.*

In the paragraph before this Lemma, Biggs and Boshier prove at least one of the integers $2mp^r - m^2q^2$ is good for the cases $m = 2$ and $m = 4$ so long as they are both positive.

The following lemma is original.

Lemma 5.4. *If $m = 4 + 8c$ for nonnegative integers c , then $2mp^r - m^2q^2$ is good if it is positive.*

Proof. Note

$$\frac{2mp^r - m^2q^2}{4} = (2 + 4c)p^r - (4 + 16c + 16c^2)q^2 \equiv 2 \pmod{4}$$

implying that $2mp^r - m^2q^2$ is good. \square

So for $m = 4, 12, 20$, $2mp^r - m^2q^2$ is good if we can show it is positive. Let r_0 be the smallest positive integer such that $p^{r_0} > 10q^2$, which makes $2mp^{r_0} - m^2q^2 = m(2p^{r_0} - mq^2)$ positive for $m < 20$. Then $p^{r_0-1} < 10q^2$ (note this is a strict inequality since p^{r_0-1} is clearly not a divisible by 10) and so

$$r_0 < 2 \log_p q + \log_p 10 + 1.$$

Thus there exists three distinct $[b] \in \Lambda(2q)$, since each value of m produces a different b_0 , which means there are three distinct non-backtracking closed walks of length $2r_0$ from the identity vertex to itself. From Lemma 3.4 we showed that this situation would imply the abelian girth of $X^{p,q}$ is at most $8r_0$. Since $n = q(q^2 + 1)$ and $d - 1 = p$ we have

$$(7) \quad \text{Abl}(X^{p,q}) \leq \frac{16}{3} \log_{d-1} n(1 + o(1))$$

for p fixed and q large.

The $16/3$ constant above may not be the optimal constant. In our arguments, we found three closed walks of with a single starting vertex and length $2r_0$. If instead we had found three walks of length r_0 sharing distinct starting and terminating vertices the argument would have given the constant of 4 in (7). Note that any closed walk from a vertex v to itself of length $2r_0$ already gives two distinct walks of length r_0 from v to the m , the middle vertex of the cycle. Identifying one more walk from v to m of length r_0 would be sufficient to improve the $16/3$ coefficient.

REFERENCES

- [ABLS07] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(04):585–603, 2007.
- [AHL02] Alon Amit, Shlomo Hoory, and Nathan Linial. A continuous analogue of the girth problem. *J. Combin. Theory Ser. B*, 84(2):340–363, 2002.
- [BB90] A. L. Biggs and N. G. Boshier. Note on the girth of ramanujan graphs. *J. Combin. Theory Ser. B*, 49(2):190–194, 1990.

- [DD89] Warren Dicks and M. J. Dunwoody. *Groups acting on graphs*. Cambridge University Press, Cambridge, 1989.
- [EL14] David Ellis and Nathan Linial. On regular hypergraphs of high girth. *The Electronic Journal of Combinatorics*, 21(1):P1–54, 2014.
- [ERS66] P Erdős, A Rényi, and VT Sós. On a problem of graph theory. *Studia Sci. Math. Hungar*, 1:215–235, 1966.
- [ES63] Paul Erdős and Horst Sachs. Reguläre graphen gegebener taillenweite mit minimaler knotenzahl. *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe*, 12:251–257, 1963.
- [FIS] Joel Friedman, Alice Izsak, and Lior Silberman. Gapped sheaves on graphs. To appear.
- [Fri15] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture: with an appendix by Warren Dicks. *Mem. Amer. Math. Soc.*, 233(1100):xii+106, 2015. With an appendix by Warren Dicks.
- [LP10] Nati Linial and Doron Puder. Word maps and spectra of random graph lifts. *Random Structures Algorithms*, 37(1):100–135, 2010.
- [LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [LS15] Roger C Lyndon and Paul E Schupp. *Combinatorial group theory*. Springer, 2015.
- [Mar88] Grigorii Aleksandrovich Margulis. Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problemy peredachi informatsii*, 24(1):51–60, 1988.
- [MN96] David JC MacKay and Radford M Neal. Near shannon limit performance of low density parity check codes. *Electronics letters*, 32(18):1645–1646, 1996.
- [TZ05] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.

DEPARTMENT OF COMPUTER SCIENCE, THE UNIVERSITY OF BRITISH COLUMBIA, 2366 MAIN MALL, V6T 1Z4, VANCOUVER BC, CANADA, AND DEPARTMENT OF MATHEMATICS, THE UNIVERSITY OF BRITISH COLUMBIA, 1984 MATHEMATICS ROAD, V6T 1Z2, VANCOUVER BC, CANADA

E-mail address: jf@cs.ubc.ca or jf@math.ubc.ca

DEPARTMENT OF COMPUTER SCIENCE, THE UNIVERSITY OF BRITISH COLUMBIA, 2366 MAIN MALL, V6T 1Z4, VANCOUVER BC, CANADA

E-mail address: aizsak@cs.ubc.ca or aizsak@gmail.com

DEPARTMENT OF MATHEMATICS, THE UNIVERSITY OF BRITISH COLUMBIA, 1984 MATHEMATICS ROAD, V6T 1Z2, VANCOUVER BC, CANADA

E-mail address: lior@math.ubc.ca

URL: <http://www.math.ubc.ca/~lior>